1 AI Service for Module 1 in the Service Life Cycle: Service

SERVICE MEISTER

IT MONITORING BY MEANS OF ADAPTIVE THRESHOLDS / SMART BASELINING

When monitoring complex IT infrastructures, many measured variables are recorded and evaluated. The monitoring is mostly done by means of static thresholds. Exceeding or falling below the limits automatically leads to the responsible administrators being notified, who then determine in a subsequent analysis whether there is actually a problem that requires remedial action. With the help of regular monitoring by means of dynamic threshold adjustment, real errors can be detected and eliminated at an early stage.

FOR THE FOLLOWING CHALLENGES
- Targeted alerting
- Faster reactions to anomalies
- Avoidance of false alarms
- Higher service availability and fewer SLA violations

THE USE CASE
The first point in monitoring is the continuous collection and storage of various measurement data (for example, performance data, data from log entries, memory usage, CPU usage, hard disk capacity, and so on).

Furthermore, in addition to visualisation, fixed threshold values are defined for comprehensive monitoring in order to be able to monitor the IT infrastructure more effectively. The correct definition of the threshold values is of great importance here.

If these thresholds are set too broadly, problems may remain undetected that may later cause a service level agreement (SLA) violation or service outage. If, on the other hand, the threshold values are set too strictly, this leads to many alarms or to manual analyses, which often reveal that there are no problems at all.

THE SOLUTION IN DETAIL
Here, the development of a statistical model for anomaly detection using dynamic thresholds directly addresses this problem of fixed thresholds.

Here, too, the measurement data is continuously recorded and processed and stored in a further step. Incorrect or incomplete measurement series are discarded.

The statistical model then compares the data in a rolling time window with measurements from the past. Only patterns that correlated with a problematic situation in the past will trigger an alarm. This leads to optimised, improved alarms.

The training of pattern recognition takes place continuously on the basis of historical data.

PROJECT STATUS
The model is under development and will be commercially available later.

REQUIREMENTS
- The service requires a trained model, for which historical data must exist.
- No special knowledge of the measured value to be recorded is required.
-

AVAILABILITY
Upon request.

CONTACT PERSON:
Henrik Oppermann
henrik.oppermann@usu.com

SERVICE MEISTER

SPECIFICATION

	Input data	Preprocessing	Data storage	Algorithms	Interfaces
High-level description	E2E measurements, measurement of server and network performance, data from log entries, memory, CPU usage, hard disk capacity	Sanity Check	NAS storage, data bases	Anomaly detection	Monitoring / Dashboard / Alarms
Configurability	Selection of the measured variables			Training period, sensitivity	Plot selection or representation of the dashboard
Technical implementation	Measurements on site	Python, Docker Container	Filesystem, DBs, Docker Container	Python, Docker Container	REST-API
Specific example from the speedboat project	The measurement data is continuously recorded	The data is processed, faulty or incomplete measurement series are discarded.	The training of pattern recognition takes place continuously on the basis of historical data.	Rolling time windows are compared with measurements from the past. An alarm is triggered only for patterns that correlate with a problematic situation in the past.	With the help of regular monitoring by means of dynamic threshold adjustment, real errors can be detected and eliminated at an early stage.

CONTACT PERSON:
Henrik Oppermann
henrik.oppermann@usu.com